

New Approximation Algorithms for Minimum Enclosing Convex Shapes

Ankan Saha* S.V.N. Vishwanathan[†] Xinhua Zhang[‡]

September 16, 2010

Abstract

Given n points in a d dimensional Euclidean space, the Minimum Enclosing Ball (MEB) problem is to find the ball with the smallest radius which contains all n points. We give a $O(ndQ/\sqrt{\epsilon})$ approximation algorithm for producing an enclosing ball whose radius is at most ϵ away from the optimum (where Q is an upper bound on the norm of the points). This improves existing results using *coresets*, which yield a $O(nd/\epsilon)$ greedy algorithm. Finding the Minimum Enclosing Convex Polytope (MECP) is a related problem wherein a convex polytope of a fixed shape is given and the aim is to find the smallest magnification of the polytope which encloses the given points. For this problem we present a $O(mndQ/\epsilon)$ approximation algorithm, where m is the number of faces of the polytope. Our algorithms borrow heavily from convex duality and recently developed techniques in non-smooth optimization, and are in contrast with existing methods which rely on geometric arguments. In particular, we specialize the excessive gap framework of [Nesterov \[2005b\]](#) to obtain our results.

*Department of Computer Science University of Chicago ankans@cs.uchicago.edu

[†]Department of Statistics and Computer Science Purdue University
vishy@stat.purdue.edu

[‡]Department of Computing Science University of Alberta xinhua.zhang.cs@gmail.com

1 Introduction

Given a set $S = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ of n points in \mathbb{R}^d , the minimum enclosing ball (MEB) is the ball with the smallest radius which contains all the points in S . The problem of finding a MEB arises in application areas as diverse as data mining, learning, statistics, computer graphics, and computational geometry [Elzinga and Hearn, 1972]. Therefore efficient algorithms for this problem are not only of theoretical interest, but also have wide practical applicability.

Exact algorithms for finding the MEB typically have an exponential dependence on d [Megiddo, 1984, Welzl, 1991]. For example, the Welzl [1991] algorithm runs in $O(n(d+1)(d+1)!)$ time which makes it inadmissible for many practical applications; in the case of linear SVMs data may have a million or more dimensions. Therefore, there has been a significant interest in finding approximation algorithms for this problem.

State of the art approximation algorithms for the MEB problem extensively use the concept of coresets [Clarkson, 2008, Badoiu and Clarkson, 2002, Panigrahy, 2004, Yildirim, 2008]. Given an $\epsilon > 0$, an ϵ -coreset $S' \subset S$ has the property that if the smallest enclosing ball containing S' is expanded by a factor of $(1+\epsilon)$, then the resulting ball also contains S . Therefore, locating an ϵ -coreset is equivalent to finding an $(1+\epsilon)$ approximation algorithm to the MEB problem. The approximation guarantees of such algorithms are **multiplicative**. Briefly, a coreset is built incrementally in a greedy fashion [Clarkson, 2008]. At every iteration, the MEB of the current candidate coreset is built. If every point in S lies in an $(1+\epsilon)$ ball of the current solution then the algorithm stops, otherwise the most *violated* point, that is, the point which is furthest away from the current MEB is included in the candidate coreset and the iterations continue. The best known algorithms in this family have a running time of $O(nd/\epsilon)$ [Panigrahy, 2004, Clarkson, 2008].

In contrast, we present a new algorithm which is derived by casting the problem of finding the MEB as a convex but non-smooth optimization problem. By specializing a general framework of Nesterov [2005b], our algorithm is able to achieve a running time of $O(ndQ/\sqrt{\epsilon})$ where Q is an upper bound on the norm of the points. Also, the approximation guarantees of our algorithm are **additive**, that is, given a tolerance $\epsilon > 0$ and denoting the optimal radius by R^* , our algorithm produces a function whose value lies between R^{*2} and $R^{*2} + \epsilon$. Although these two types of approximation guarantees seem different, by a simple argument in section 3.2, we show that our algorithm also yields a traditional scale-invariant ϵ multiplicative approximation with $O(ndQ/\sqrt{\epsilon})$ effort.

We extend our analysis to the closely related minimum enclosing convex polytope (MECP) problem, and present a new algorithm. As before, given a set $S = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ of n points in \mathbb{R}^d , the task here is to find the smallest polytope of a given fixed shape which encloses these points. In our setting translations and magnifications are allowed but rotations are not allowed. We present a $O(mndQ/\epsilon)$ approximation algorithm, where m denotes the number of faces of the polytope.

We apply our algorithms to two problems of interest in machine learning namely finding the maximum margin hyperplane and computing the distance of a polytope from the origin. A coresets algorithm for the first problem was proposed by Har-Peled et al. [2007] while the second one was studied by Gärtner and Jaggi [2009]. In both cases our algorithms require fewer number of iterations and yield better computational complexity bounds.

Our paper is structured as follows: In Section 2 we introduce notation, briefly review some results from convex duality, and present the general framework of Nesterov [2005b]. In Section 3 we address the MEB problem and in Section 4 the MECP problem, and present our algorithms and their analysis. We discuss some applications of our results to machine learning problems in Section 5. The paper then concludes with a discussion and outlook for the future in Section 6. Technical proofs can be found in Appendix A and B, while preliminary experimental evaluation can be found in Appendix D.

2 Definitions and Preliminaries

In this paper, lower bold case letters (*e.g.*, $\mathbf{w}, \boldsymbol{\mu}$) denote vectors, while upper bold case letters (*e.g.*, \mathbf{A}) denote matrices or linear operators. We use w_i to denote the i -th component of \mathbf{w} , A_{ij} to denote the (i, j) -th entry of \mathbf{A} , and $\langle \mathbf{w}, \mathbf{w}' \rangle := \sum_i w_i w'_i$ to denote the Euclidean dot product between vectors \mathbf{w} and \mathbf{w}' . Δ_k denotes the k dimensional simplex. Unless specified otherwise, $\|\cdot\|$ refers to the Euclidean norm $\|\mathbf{w}\| := \sqrt{\langle \mathbf{w}, \mathbf{w} \rangle} = (\sum_{i=1}^n w_i^2)^{\frac{1}{2}}$. For a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, we have the following definition of the norm

$$\|\mathbf{A}\| = \max \{ \langle \mathbf{A}\mathbf{w}, \mathbf{u} \rangle : \|\mathbf{w}\| = 1, \|\mathbf{u}\| = 1 \}.$$

We also denote $\overline{\mathbb{R}} := \mathbb{R} \cup \{\infty\}$, and $[t] := \{1, \dots, t\}$.

Definition 1 Let $Q_1 \subseteq \mathbb{R}^n$, $f : Q_1 \rightarrow \overline{\mathbb{R}}$, and $f^* := \min_{\mathbf{w} \in Q_1} f(\mathbf{w}) < \infty$. A point $\mathbf{w}' \in Q_1$ such that

$$f(\mathbf{w}') \leq f^* + \epsilon \tag{1}$$

is said to be an ϵ -accurate minimizer of f . We will also sometimes call \mathbf{w}' an ϵ -accurate solution.

The following three standard concepts from convex analysis (see *e.g.* [Hiriart-Urruty and Lemaréchal \[1993\]](#)) are extensively used in the sequel.

Definition 2 A convex function $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is strongly convex with respect to a norm $\|\cdot\|$ if there exists a constant $\rho > 0$ such that $f - \frac{\rho}{2}\|\cdot\|^2$ is convex. ρ is called the modulus of strong convexity of f , and for brevity we will call f ρ -strongly convex or ρ -s.c.

Definition 3 Suppose a function $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is differentiable on $Q \subseteq \mathbb{R}^n$. Then f is said to have Lipschitz continuous gradient (l.c.g) with respect to a norm $\|\cdot\|$ if there exists a constant L such that

$$\|\nabla f(\mathbf{w}) - \nabla f(\mathbf{w}')\| \leq L\|\mathbf{w} - \mathbf{w}'\| \quad \forall \mathbf{w}, \mathbf{w}' \in Q. \quad (2)$$

For brevity, we will call f L -l.c.g.

Definition 4 The Fenchel dual of a function $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is a function $f^* : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ defined by

$$f^*(\mathbf{w}^*) = \sup_{\mathbf{w} \in \mathbb{R}^n} \{\langle \mathbf{w}, \mathbf{w}^* \rangle - f(\mathbf{w})\} \quad (3)$$

Strong convexity and Lipschitz continuity of the gradient are related by Fenchel duality according to the following lemma:

Lemma 5 ([Hiriart-Urruty and Lemaréchal \[1993, Theorem 4.2.1 and 4.2.2\]](#))

1. If $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is ρ -s.c., then f^* is finite on \mathbb{R}^n and f^* is $\frac{1}{\rho}$ -l.c.g.
2. If $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex, differentiable on \mathbb{R}^n , and L -l.c.g, then f^* is $\frac{1}{L}$ -s.c.

2.1 Nesterov's Framework

In sections 3 and 4 we will show that the MEB and MECP problems respectively can be cast as minimizing convex non-smooth objective functions. In a series of papers, Nesterov [[Nesterov, 1983, 2005a,b](#)] proposed a general framework for this task, which we now briefly review.

Let Q_1 and Q_2 be subsets of Euclidean spaces and \mathbf{A} be a linear map from Q_1 to Q_2 . Suppose f and g are convex functions defined on Q_1 and Q_2 respectively, and we are interested in the following optimization problem:

$$\min_{\mathbf{w} \in Q_1} J(\mathbf{w}) \text{ where } J(\mathbf{w}) := f(\mathbf{w}) + g^*(\mathbf{A}\mathbf{w}) = f(\mathbf{w}) + \max_{\mathbf{u} \in Q_2} \{\langle \mathbf{A}\mathbf{w}, \mathbf{u} \rangle - g(\mathbf{u})\}. \quad (4)$$

We will make the following standard assumptions: a) Q_2 is compact; b) with respect to a certain norm on Q_1 , the function f defined on Q_1 is ρ -s.c. but not necessarily $l.c.g$, and c) with respect to a certain norm on Q_2 , the function g defined on Q_2 is L_g - $l.c.g$ and convex, but not necessarily strongly convex.

The key difficulty in solving (4) arises because g^* and hence J may be non-smooth. Our aim is to uniformly approximate $J(\mathbf{w})$ with a smooth and strongly convex function. Towards this end let d be a σ -s.c. smooth function with the following properties:

$$\min_{\mathbf{u} \in Q_2} d(\mathbf{u}) = 0, \quad \mathbf{u}_0 = \operatorname{argmin}_{\mathbf{u} \in Q_2} d(\mathbf{u}), \text{ and } \mathcal{D} := \max_{\mathbf{u} \in Q_2} d(\mathbf{u}).$$

In optimization parlance d is called a prox-function. For a positive constant $\mu \in \mathbb{R}$ define

$$J_\mu(\mathbf{w}) := f(\mathbf{w}) + \max_{\mathbf{u} \in Q_2} \{\langle \mathbf{A}\mathbf{w}, \mathbf{u} \rangle - g(\mathbf{u}) - \mu d(\mathbf{u})\}. \quad (5)$$

It can be easily verified that J_μ is not only smooth and convex but also $\frac{1}{\sigma\mu}\|\mathbf{A}\|^2$ - $l.c.g$ [Nesterov, 2005a]. Furthermore, if $\mathcal{D} < \infty$ then J_μ is uniformly close to J , that is,

$$J_\mu(\mathbf{w}) \leq J(\mathbf{w}) \leq J_\mu(\mathbf{w}) + \mu \mathcal{D}. \quad (6)$$

If some mild constraint qualifications hold [*e.g.* Theorem 3.3.5 Borwein and Lewis, 2000] one can write the dual $D(\mathbf{u})$ of $J(\mathbf{w})$ using \mathbf{A}^\top (the transpose of \mathbf{A}) as

$$D(\mathbf{u}) := -g(\mathbf{u}) - f^*(-\mathbf{A}^\top \mathbf{u}) = -g(\mathbf{u}) - \max_{\mathbf{w} \in Q_1} \{\langle -\mathbf{A}\mathbf{w}, \mathbf{u} \rangle - f(\mathbf{w})\}, \quad (7)$$

and assert the following:

$$\inf_{\mathbf{w} \in Q_1} J(\mathbf{w}) = \sup_{\mathbf{u} \in Q_2} D(\mathbf{u}), \quad \text{and} \quad J(\mathbf{w}) \geq D(\mathbf{u}) \quad \forall \mathbf{w} \in Q_1, \mathbf{u} \in Q_2. \quad (8)$$

The key idea of excessive gap minimization pioneered by [Nesterov \[2005b\]](#) is to maintain two estimation sequences $\{\mathbf{w}_k\}$ and $\{\mathbf{u}_k\}$, together with a diminishing sequence $\{\mu_k\}$ such that

$$\boxed{J_{\mu_k}(\mathbf{w}_k) \leq D(\mathbf{u}_k), \text{ and } \lim_{k \rightarrow \infty} \mu_k = 0.} \quad (9)$$

The idea is illustrated in Figure 1. In conjunction with (8) and (6), it is not hard to see that $\{\mathbf{w}_k\}$ and $\{\mathbf{u}_k\}$ approach the solution of $\min_{\mathbf{w}} J(\mathbf{w}) = \max_{\mathbf{u}} D(\mathbf{u})$. Using (6), (5), and (9), we can derive the following bound on the duality gap:

$$J(\mathbf{w}_k) - D(\mathbf{u}_k) \leq J_{\mu_k}(\mathbf{w}_k) + \mu_k \mathcal{D} - D(\mathbf{u}_k) \leq \mu_k \mathcal{D}. \quad (10)$$

In other words, the duality gap is reduced at the same rate at which μ_k approaches 0. To turn this idea into an implementable algorithm we need to answer the following two questions:

1. How to efficiently find initial points \mathbf{w}_1 , \mathbf{u}_1 and μ_1 that satisfy (9).
2. Given \mathbf{w}_k , \mathbf{u}_k , and μ_k , how to *efficiently* find iterates \mathbf{w}_{k+1} , \mathbf{u}_{k+1} , and μ_{k+1} which maintain (9). To achieve the best possible convergence rate it is desirable to anneal μ_k as fast as possible while still allowing \mathbf{w}_k and \mathbf{u}_k to be updated efficiently.

We will now show how the MEB and MECP problems can be cast as convex optimization problems and derive implementable algorithms by answering the above questions.

3 Minimum Enclosing Ball

Given a set of n points $S = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ in a d dimensional space \mathbb{R}^d , a Euclidean ball $B(\mathbf{c}, R)$ of radius R centered at \mathbf{c} is said to be an enclosing ball if $\forall i \in [n], \mathbf{x}_i \in B(\mathbf{c}, R)$.

3.1 Formulation as an Optimization Problem

Clearly, $\mathbf{x}_i \in B(\mathbf{c}, R)$ if, and only if, $\|\mathbf{c} - \mathbf{x}_i\|^2 \leq R^2$. Using this observation, the MEB problem can be cast as the following optimization problem:

$$\min_{R \in \mathbb{R}} R \quad \text{s.t.} \quad \|\mathbf{c} - \mathbf{x}_i\|^2 \leq R^2 \quad \forall i,$$




Figure 1: Illustration of excessive gap. By strong convexity, the dual $D(\mathbf{u})$ is always a lower bound to the primal $J(\mathbf{w})$ (left). We approximate $J(\mathbf{w})$ by a smooth lower bound $J_{\mu_k}(\mathbf{w})$ (middle). Since $D(\mathbf{u}_k)$ is sandwiched between $J_{\mu_k}(\mathbf{w})$ and $J(\mathbf{w})$, as $\mu_k \rightarrow 0$ we get closer and closer to the true optimum (right).

which in turn can be reformulated as

$$\min_{\mathbf{c} \in \mathbb{R}^d} \max_{\mathbf{x}_i \in S} \|\mathbf{c} - \mathbf{x}_i\|^2. \quad (11)$$

Rearranging terms

$$\min_{\mathbf{c} \in \mathbb{R}^d} J(\mathbf{c}) = \|\mathbf{c}\|^2 + \max_{\mathbf{x}_i \in S} \{-2\langle \mathbf{c}, \mathbf{x}_i \rangle + \|\mathbf{x}_i\|^2\} = \|\mathbf{c}\|^2 + \max_{\mathbf{u} \in \Delta_n} \{\langle \mathbf{A}\mathbf{c}, \mathbf{u} \rangle + \langle \mathbf{u}, \mathbf{b} \rangle\}, \quad (12)$$

where $\mathbf{A} = -2[\mathbf{x}_1, \mathbf{x}_2 \dots \mathbf{x}_n]^\top$, and $b_i = \|\mathbf{x}_i\|^2$. Clearly, $J(\mathbf{c})$ can be identified with (4) by setting $Q_1 = \mathbb{R}^d$, $Q_2 = \Delta_n$, $f(\mathbf{c}) = \|\mathbf{c}\|^2$, and $g(\mathbf{u}) = -\langle \mathbf{u}, \mathbf{b} \rangle$. It can be verified that g is 0-*l.c.g.*, while f is 2-*s.c.* Therefore, one can employ Nesterov's framework (Section 2.1) to minimize $J(\mathbf{c})$. However, as we stated before, we need to specialize the framework to our setting to obtain an efficient and implementable algorithm. Towards this end note that the Fenchel dual of $\|\cdot\|^2$ is $\frac{1}{4} \|\cdot\|^2$, and use (7) to write the dual of (12) as

$$D(\mathbf{u}) = \langle \mathbf{u}, \mathbf{b} \rangle - \frac{1}{4} \mathbf{u}^\top \mathbf{A} \mathbf{A}^\top \mathbf{u}. \quad (13)$$

By using the Cauchy-Schwartz inequality, the gradient

$$\nabla D(\mathbf{u}) = \mathbf{b} - \frac{1}{2} \mathbf{A} \mathbf{A}^\top \mathbf{u}, \quad (14)$$

can be shown to satisfy

$$\|\nabla D(\mathbf{u}_1) - \nabla D(\mathbf{u}_2)\| = \left\| -\frac{1}{2} \mathbf{A} \mathbf{A}^\top \mathbf{u}_1 + \frac{1}{2} \mathbf{A} \mathbf{A}^\top \mathbf{u}_2 \right\| \leq \frac{1}{2} \|\mathbf{A} \mathbf{A}^\top\| \|\mathbf{u}_1 - \mathbf{u}_2\|, \quad (15)$$

thus establishing that $D(\mathbf{u})$ is $\frac{1}{2} \|\mathbf{A}\mathbf{A}^\top\|$ -l.c.g. We define $L = \frac{1}{2} \|\mathbf{A}\mathbf{A}^\top\|$.

Next we turn our attention to the prox-function. Recall that Nesterov's framework requires a σ -s.c. prox-function on $Q_2 = \Delta_n$; in our case we set

$$d(\mathbf{u}) = \frac{\sigma}{2} \|\mathbf{u} - \mathbf{u}_0\|^2$$

where $\mathbf{u}_0 = (\frac{1}{n}, \dots, \frac{1}{n}) \in \mathbb{R}^n$. For this choice $\mathbf{u}_0 = \operatorname{argmin}_{\mathbf{u} \in \Delta_n} d(\mathbf{u})$, $d(\mathbf{u}_0) = 0$, and

$$\mathcal{D} = \max_{\mathbf{u} \in \Delta_n} d(\mathbf{u}) = \frac{\sigma}{2} \max_{\mathbf{u} \in \Delta_n} \|\mathbf{u} - \mathbf{u}_0\|^2 = \frac{\sigma}{2} \left(1 - \frac{1}{n}\right)^2 \leq \frac{\sigma}{2}. \quad (16)$$

Furthermore, for notational convenience, define the following three maps:

$$\mathbf{c}(\mathbf{u}) = \operatorname{argmin}_{\mathbf{c} \in \mathbb{R}^d} \left\{ \langle \mathbf{A}\mathbf{c}, \mathbf{u} \rangle + \|\mathbf{c}\|^2 \right\} \quad (17)$$

$$\mathbf{u}_\mu(\mathbf{c}) = \operatorname{argmax}_{\mathbf{u} \in \Delta_n} \left\{ \langle \mathbf{A}\mathbf{c}, \mathbf{u} \rangle + \langle \mathbf{u}, \mathbf{b} \rangle - \mu d(\mathbf{u}) \right\} = \operatorname{argmin}_{\mathbf{u} \in \Delta_n} \left\{ \frac{\mu\sigma}{2} \|\mathbf{u} - \mathbf{u}_0\|^2 - \langle \mathbf{A}\mathbf{c} + \mathbf{b}, \mathbf{u} \rangle \right\}. \quad (18)$$

$$V(\mathbf{u}) = \operatorname{argmin}_{\mathbf{v} \in \Delta_n} \left\{ \frac{L}{2} \|\mathbf{v} - \mathbf{u}\|^2 - \langle \nabla D(\mathbf{u}), \mathbf{v} - \mathbf{u} \rangle \right\} = \operatorname{argmin}_{\mathbf{v} \in \Delta_n} \left\{ \frac{L}{2} \|\mathbf{v} - \mathbf{u}\|^2 - \left\langle \mathbf{b} - \frac{1}{2} \mathbf{A}\mathbf{A}^\top \mathbf{u}, \mathbf{v} - \mathbf{u} \right\rangle \right\}. \quad (19)$$

With this notation in place we now describe our excessive gap minimization method in Algorithm 1. Unrolling the recursive update for μ_k yields

$$\mu_k = (1 - \tau_{k-1}) \mu_{k-1} = \frac{k}{k+2} \mu_{k-1} = \frac{(k)(k-1) \dots 2}{(k+2)(k+1) \dots 4} \frac{L}{\sigma} = \frac{6}{(k+1)(k+2)} \frac{L}{\sigma}. \quad (20)$$

Plugging this into (10) and using (16) immediately yields the following theorem:

Theorem 6 (Duality gap) *The sequences $\{\mathbf{c}_k\}$ and $\{\mathbf{u}_k\}$ in Algorithm 1 satisfy*

$$J(\mathbf{c}_k) - D(\mathbf{u}_k) \leq \frac{6LD}{\sigma(k+1)(k+2)} \leq \frac{3L}{(k+1)(k+2)}. \quad (21)$$

As is standard [see e.g. Yildirim, 2008], if we assume that the input data points lie inside a ball of radius \mathcal{Q} , that is, $\max_i \|\mathbf{x}_i\| \leq \mathcal{Q}$ then we can write

$$L = \frac{1}{2} \|\mathbf{A}\mathbf{A}^\top\| = \frac{1}{2} \max_{\|\mathbf{c}\|=\|\mathbf{u}\|=1} \mathbf{c}^\top \mathbf{A}\mathbf{A}^\top \mathbf{u} = \frac{1}{2} \max_{\|\mathbf{u}\|=1} \left\| \mathbf{A}^\top \mathbf{u} \right\|^2 = 2 \max_i \|\mathbf{x}_i\|^2 = 2\mathcal{Q}^2. \quad (22)$$

Algorithm 1: Excessive gap minimization applied to MEB

Output: Sequences $\{\mathbf{c}_k\}$, $\{\mathbf{u}_k\}$, and $\{\mu_k\}$ that satisfy (9), with $\lim_{k \rightarrow \infty} \mu_k = 0$.

```

1 Initialize: Let  $\mathbf{u}_0 = (\frac{1}{n}, \dots, \frac{1}{n})$ ,  $\mu_1 = \frac{L}{\sigma}$ ,  $\mathbf{c}_1 = \mathbf{c}(\mathbf{u}_0)$ ,  $\mathbf{u}_1 = V(\mathbf{u}_0)$ .;
2 for  $k = 1, 2, \dots$  do
3    $\tau_k \leftarrow \frac{2}{k+3}$ .
4    $\beta_k \leftarrow (1 - \tau_k)\mathbf{u}_k + \tau_k \mathbf{u}_{\mu_k}(\mathbf{c}_k)$ .
5    $\mathbf{c}_{k+1} \leftarrow (1 - \tau_k)\mathbf{c}_k + \tau_k \mathbf{c}(\beta_k)$ .
6    $\mathbf{u}_{k+1} \leftarrow V(\beta_k)$ .
7    $\mu_{k+1} \leftarrow (1 - \tau_k)\mu_k$ .

```

The last equality follows because $\mathbf{A} = -2[\mathbf{x}_1, \mathbf{x}_2 \dots \mathbf{x}_n]^\top$ and the maximum is attained by setting $\mathbf{u} = \mathbf{e}_j$ where $j = \operatorname{argmax}_i \|\mathbf{x}_i\|^2$. Note that this is just a conservative estimate and a larger value of L also guarantees convergence of the algorithm. Plugging this back into (21) yields

$$J(\mathbf{c}_k) - D(\mathbf{u}_k) \leq \frac{6\mathcal{Q}^2}{(k+1)(k+2)}. \quad (23)$$

Therefore, to obtain an ϵ accurate solution of (12) it suffices to ensure that

$$\frac{6\mathcal{Q}^2}{(k+1)(k+2)} \leq \epsilon. \quad (24)$$

Solving for k yields the $O(\mathcal{Q}/\sqrt{\epsilon})$ bounds on the number of iterations as claimed. All that remains is to show that

Theorem 7 *The update rule of Algorithm 1 guarantees that (9) is satisfied for all $k \geq 1$.*

Proof See Appendix A. ■

Each iteration of Algorithm 1 requires us to compute $\mathbf{c}(\mathbf{u})$, $\mathbf{u}_\mu(\mathbf{c})$, and $V(\mathbf{u})$ (see (17), (18), and (19)). All other operations either require constant or linear time. We now show that each of these three maps can be computed in $O(nd)$ time. This in conjunction with Theorem 6 shows that the time complexity of our algorithm to find an ϵ accurate solution of (12) is $O(nd/\sqrt{\epsilon})$.

By computing the gradient of $\langle \mathbf{A}\mathbf{c}, \mathbf{u} \rangle + \|\mathbf{c}\|^2$ and setting it to zero we can show that

$$\mathbf{c}(\mathbf{u}) = -\frac{1}{2}\mathbf{A}^\top \mathbf{u}. \quad (25)$$

Since \mathbf{A} is a $n \times d$ matrix computing $\mathbf{c}(\mathbf{u})$ takes $O(nd)$ time. On the other hand, computation of $\mathbf{u}_\mu(\mathbf{c})$ can be cast as the following Quadratic programming (QP) problem with linear constraints:

$$\begin{aligned} \min_{\mathbf{u}} \quad & \frac{\mu\sigma}{2} \|\mathbf{u}\|^2 - \langle \mathbf{A}\mathbf{c} + \mathbf{b} + \mu\sigma\mathbf{u}_0, \mathbf{u} \rangle \\ \text{s.t.} \quad & \sum_i u_i = 1 \text{ and } 0 \leq u_i \leq 1. \end{aligned} \quad (26)$$

Computing $\mathbf{A}\mathbf{c} + \mathbf{b} + \mu\sigma\mathbf{u}_0$ requires $O(nd)$ time. Given $\mathbf{A}\mathbf{c} + \mathbf{b} + \mu\sigma\mathbf{u}_0$, we show in Appendix B that the above QP can be solved in $O(n)$ time. Finally, after some simple algebraic manipulation computation of $V(\mathbf{u})$ can be also be cast as a Quadratic programming (QP) problem with linear constraints as follows:

$$\begin{aligned} \min_{\mathbf{v}} \quad & \frac{L}{2} \|\mathbf{v}\|^2 - \left\langle L\mathbf{u} - \frac{1}{2}\mathbf{A}\mathbf{A}^\top \mathbf{u} + \mathbf{b}, \mathbf{v} \right\rangle \\ \text{s.t.} \quad & \sum_i v_i = 1 \text{ and } 0 \leq v_i \leq 1. \end{aligned} \quad (27)$$

Again, the computational bottleneck is in computing $L\mathbf{u} - \frac{1}{2}\mathbf{A}\mathbf{A}^\top \mathbf{u} + \mathbf{b}$ which takes $O(nd)$ effort^a. Given $L\mathbf{u} - \frac{1}{2}\mathbf{A}\mathbf{A}^\top \mathbf{u} + \mathbf{b}$ the algorithm in Appendix B can be applied to solve (27) in $O(n)$ time.

3.2 Multiplicative versus Additive Approximation: Scale Invariance

Existing approximation algorithms for the MEB problem based on coresets provide a multiplicative approximation. Given a set of points, the coreset algorithms output a center \mathbf{c} and radius R such that all the given points lie inside the ball of radius $R(1 + \epsilon)$ centered at \mathbf{c} . In contrast, our algorithm produces a center and a radius R such that $R^{*2} \leq R^2 \leq R^{*2} + \epsilon'$, where R^* denotes the radius of the optimal minimum enclosing ball.

At first glance the two types of guarantees do not look directly comparable since the additive guarantees seem to vary with change of scale. To produce an ϵ_M *scale invariant* multiplicative approximation with our algorithm,

^aTo compute $\mathbf{A}\mathbf{A}^\top \mathbf{u}$ efficiently we first compute $\mathbf{a} = \mathbf{A}^\top \mathbf{u}$ and then compute $\mathbf{A}\mathbf{a}$.

set $\epsilon' = \epsilon_M R^{*2}$. In view of (24) it follows that $R^{*2} \leq R^2 \leq R^{*2}(1 + \epsilon_M)$ whenever^b

$$\frac{6\mathcal{Q}^2}{(k+1)(k+2)} \leq R^{*2}\epsilon_M. \quad (28)$$

Solving for k obtains

$$\frac{\mathcal{Q}}{R^*} \sqrt{\frac{6}{\epsilon_M}} \leq k. \quad (29)$$

However, since R^* is unknown this bound cannot be used as a practical stopping criterion. Instead, one can use the following observation: select an arbitrary pair of points \mathbf{x}_i and \mathbf{x}_j from S and compute $\frac{1}{2} \|\mathbf{x}_i - \mathbf{x}_j\|$. Denote this quantity by \mathcal{P} , and let \mathbf{c}^* be the center of the optimal MEB. Clearly

$$\mathcal{P} = \frac{1}{2} \|\mathbf{x}_i - \mathbf{x}_j\| \leq \frac{1}{2} \|\mathbf{x}_i - \mathbf{c}^*\| + \frac{1}{2} \|\mathbf{x}_j - \mathbf{c}^*\| = \frac{1}{2}R^* + \frac{1}{2}R^* = R^*. \quad (30)$$

Therefore replacing R^* by \mathcal{P} in (29) yields the following computable upper bound on the number of iterations:

$$\frac{\mathcal{Q}}{\mathcal{P}} \sqrt{\frac{6}{\epsilon_M}} \leq k. \quad (31)$$

This shows that $O(1/\sqrt{\epsilon_M})$ iterations of our algorithm suffice to produce a ϵ_M -multiplicative approximation.

Note that just like in the case of coresets based algorithms this bound is scaling invariant, that is, if all the \mathbf{x}_i in S are scaled by a factor $\alpha > 0$ the bound still holds. To see this one merely has to observe that after scaling \mathcal{P} becomes $\alpha\mathcal{P}$ and \mathcal{Q} becomes $\alpha\mathcal{Q}$, but the ratio \mathcal{Q}/\mathcal{P} which appears in (31) remains unchanged. Thus while our algorithm is modeled as a generic optimization formulation, it is possible to obtain scale invariance by appropriately choosing ϵ' to be $\epsilon_M \mathcal{P}^2$.

4 Minimum Enclosing Convex Polytope

In the Minimum Enclosing Convex Polytope (MECP) problem we are given a polytope of fixed shape which can be translated and magnified but rotations are not allowed. Furthermore, we assume that the polytope has a finite

^bWe prove our bounds in terms of R^2 but it is trivial to convert this to a bound in terms of R .

number of faces and hence it can be expressed as an intersection of a finite number of hyperplanes:

$$\langle \mathbf{w}_i, \mathbf{x} - \mathbf{c} \rangle \leq t_i \quad i = 1, 2, \dots, m,$$

where \mathbf{c} is the *center* of the polytope about which it is magnified. Given a set of points $S = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ we want to find the minimum magnification of the convex polytope that encloses all the points in S .

4.1 Formulation as an Optimization problem

Clearly an enclosing polytope is one for which $\langle \mathbf{w}_i, \mathbf{x}_j - \mathbf{c} \rangle \leq t_i$ for all i and j . This observation helps us to cast the MECP problem as the following optimization problem

$$\min_{R \in \mathbb{R}, \mathbf{c} \in Q_1} R \quad \text{s.t.} \quad \langle \mathbf{w}_i, \mathbf{x}_j - \mathbf{c} \rangle \leq R t_i$$

or equivalently as

$$\min_{R \in \mathbb{R}, \mathbf{c} \in Q_1} R \quad \text{s.t.} \quad \left\langle \frac{\mathbf{w}_i}{t_i}, \mathbf{x}_j - \mathbf{c} \right\rangle \leq R \quad \forall i, j.$$

Here R is the scale of magnification of the polygon and $Q_1 \subset \mathbb{R}^d$ is a bounded set, for example, a ball of certain fixed radius \mathcal{Q} centered at the origin which is assumed to contain the solution \mathbf{c} . Usually Q_1 is taken to be a ball which contains all the points in S but this need not always be the case. Also we will assume that \mathbf{w}_i/t_i lie inside a ball of radius \mathcal{W} . Writing $\tilde{\mathbf{w}}_i = \mathbf{w}_i/t_i$ the problem can be expressed as,

$$\min_{\mathbf{c} \in Q_1} \max_{i,j} \langle \tilde{\mathbf{w}}_i, \mathbf{x}_j - \mathbf{c} \rangle = \min_{\mathbf{c} \in Q_1} \max_{\mathbf{u} \in \Delta_{mn}} \sum_{i,j} u_{ij} \langle \tilde{\mathbf{w}}_i, \mathbf{x}_j - \mathbf{c} \rangle. \quad (32)$$

Notice that, as before, we have replaced the maximization over a finite set by a maximization over the simplex. Clearly this problem can be rewritten as

$$\min_{\mathbf{c} \in Q_1} J(\mathbf{c}) \text{ where } J(\mathbf{c}) = \max_{\mathbf{u} \in \Delta_{mn}} \sum_{i,j} u_{ij} \langle \tilde{\mathbf{w}}_i, \mathbf{x}_j - \mathbf{c} \rangle = \max_{\mathbf{u} \in \Delta_{mn}} \{ \langle \mathbf{A}\mathbf{c}, \mathbf{u} \rangle + \langle \mathbf{u}, \mathbf{b} \rangle \}. \quad (33)$$

We used $\mathbf{A} = [\underbrace{-\tilde{\mathbf{w}}_1, \dots}_{n \text{ times}} \underbrace{-\tilde{\mathbf{w}}_2, \dots, \dots}_{n \text{ times}} - \tilde{\mathbf{w}}_m]^\top$, a $mn \times d$ matrix with each $\tilde{\mathbf{w}}_i$ repeated n times as the columns and \mathbf{b} a mn dimensional vector with

$b_{ij} = \langle \mathbf{w}_i, \mathbf{x}_j \rangle$ to write the above expression. Clearly $J(\mathbf{c})$ can be identified with (4) by setting $Q_2 = \Delta_{mn}$, $f(\mathbf{c}) = 0$ and $g(\mathbf{u}) = -\langle \mathbf{b}, \mathbf{u} \rangle$. Here, g is 0-*l.c.g*, however, f is no longer strongly convex. Therefore, we will work with the following function

$$J_\eta(\mathbf{c}) = \eta \|\mathbf{c}\|^2 + \max_{\mathbf{u} \in \Delta_{mn}} \{ \langle \mathbf{A}\mathbf{c}, \mathbf{u} \rangle - g(\mathbf{u}) \}.$$

Let $J_\eta^* = \min_{\mathbf{c} \in Q_1} J_\eta(\mathbf{c})$ and $J^* = \min_{\mathbf{c} \in Q_1} J(\mathbf{c})$. Since $\|\mathbf{c}\|^2 \leq Q^2$ for all $\mathbf{c} \in Q_1$ we have that

$$J_\eta^* \leq J^* + \eta Q^2.$$

Suppose we can minimize J_η to $\epsilon/2$ precision, that is, find a \mathbf{c} such that $J_\eta(\mathbf{c}) \leq J_\eta^* + \epsilon/2$ then the above observation allows us to write the following series of inequalities

$$J(\mathbf{c}) \leq J_\eta(\mathbf{c}) \leq J_\eta^* + \frac{\epsilon}{2} \leq J^* + \frac{\epsilon}{2} + \eta Q^2.$$

In other words, every $\epsilon/2$ accurate solution of J_η is a $\epsilon/2 + \eta Q^2$ accurate solution of J . In particular, if we set $\eta = \epsilon/2Q^2$ then every $\epsilon/2$ accurate solution of J_η is an ϵ accurate solution of J . Furthermore, J_η is nearly identical to (12) except that $\|\mathbf{c}\|^2$ is now replaced by $\frac{\epsilon}{2Q^2} \|\mathbf{c}\|^2$. Consequently, Algorithm 1 can be directly applied to minimize J_η with the following changes:

$$D_\eta(\mathbf{u}) = \langle \mathbf{u}, \mathbf{b} \rangle - \frac{Q^2}{2\epsilon} \mathbf{u}^\top \mathbf{A} \mathbf{A}^\top \mathbf{u}. \quad (34)$$

$$\nabla D_\eta(\mathbf{u}) = \mathbf{b} - \frac{Q^2}{\epsilon} \mathbf{A} \mathbf{A}^\top \mathbf{u}. \quad (35)$$

$$\begin{aligned} V(\mathbf{u}) &= \operatorname{argmin}_{\mathbf{v} \in \Delta_{mn}} \left\{ \frac{L}{2} \|\mathbf{v} - \mathbf{u}\|^2 - \langle \nabla D_\eta(\mathbf{u}), \mathbf{v} - \mathbf{u} \rangle \right\} \\ &= \operatorname{argmin}_{\mathbf{v} \in \Delta_{mn}} \left\{ \frac{L}{2} \|\mathbf{v} - \mathbf{u}\|^2 - \left\langle \mathbf{b} - \frac{Q^2}{\epsilon} \mathbf{A} \mathbf{A}^\top \mathbf{u}, \mathbf{v} - \mathbf{u} \right\rangle \right\}. \end{aligned} \quad (36)$$

A simple application of the Cauchy-Schwartz inequality shows that

$$\|\nabla D_\eta(\mathbf{u}_1) - \nabla D_\eta(\mathbf{u}_2)\| = \left\| -\frac{Q^2}{\epsilon} \mathbf{A} \mathbf{A}^\top \mathbf{u}_1 + \frac{Q^2}{\epsilon} \mathbf{A} \mathbf{A}^\top \mathbf{u}_2 \right\| \leq \frac{Q^2}{\epsilon} \|\mathbf{A} \mathbf{A}^\top\| \|\mathbf{u}_1 - \mathbf{u}_2\|, \quad (37)$$

thus establishing that $D_\eta(\mathbf{u})$ is $\frac{Q^2}{\epsilon} \|\mathbf{A} \mathbf{A}^\top\|$ -*l.c.g*. We define $L_\eta = \frac{Q^2}{\epsilon} \|\mathbf{A} \mathbf{A}^\top\|$. Since $\tilde{\mathbf{w}}_i$ are assumed to lie inside a ball of radius \mathcal{W} , by an argument analogous to (22), it follows that $L_\eta = Q^2 \mathcal{W}^2 / \epsilon$. Plugging this into (21) obtains

$$J(\mathbf{c}_k) - D(\mathbf{u}_k) \leq \frac{3Q^2 \mathcal{W}^2}{\epsilon(k+1)(k+2)}. \quad (38)$$

In order to obtain an $\epsilon/2$ accurate solution we need to solve for k by setting $\frac{3\mathcal{Q}^2\mathcal{W}^2}{\epsilon(k+1)(k+2)} \leq \epsilon/2$. This yields $k \geq \sqrt{6}\mathcal{Q}\mathcal{W}/\epsilon$, which shows the $O(\mathcal{Q}/\epsilon)$ iteration bound as claimed. Extending the arguments for MEB to the MECF case, the per iteration complexity of $O(mnd)$ can readily be established. We omit details for brevity.

5 Applications to Machine Learning

The connection between the MEB problem and SVMs has been discussed in a number of publications [Clarkson, 2008, Har-Peled et al., 2007, Gärtner and Jaggi, 2009]. Practical algorithms using coresets were also proposed in Tsang et al. [2007] and Tsang et al. [2005]. In all these cases, our improved MEB algorithm can be plugged in as a subroutine and will yield corresponding speedups. We describe these kernel based algorithms and their connection with MEB in appendix C. In this section, we present two machine learning problems wherein our algorithms lead to better bounds than existing coreset based approaches.

5.1 Finding Large Margin Classifiers

In Har-Peled et al. [2007] a coreset algorithm (Coreset SVM) for finding the maximum margin hyperplane was described. It turns out that our MECF algorithm can be specialized to their setting, and yields improved bounds. Briefly, given m labeled data points (\mathbf{z}_i, y_i) with $\mathbf{z}_i \in \mathbb{R}^d$ and $y_i \in \{\pm 1\}$ the maximum margin hyperplane can be found by solving^c $\operatorname{argmax}_{\mathbf{c} \in \mathbb{R}^d, \|\mathbf{c}\|=1} \min_i y_i \langle \mathbf{z}_i, \mathbf{c} \rangle$. Equivalently, by defining $\tilde{\mathbf{w}}_i = y_i \cdot \mathbf{z}_i$ one can solve

$$\operatorname{argmin}_{\mathbf{c} \in \mathbb{R}^d, \|\mathbf{c}\|=1} \max_i \langle \tilde{\mathbf{w}}_i, -\mathbf{c} \rangle. \quad (39)$$

The above problem can be identified with (32) if we set S to be the empty set and Q_1 to be $\{\mathbf{c} \in \mathbb{R}^d \text{ s.t. } \|\mathbf{c}\| = 1\}$. With this substitution our MECF algorithm can directly be applied to solve (39). In this case $\mathcal{Q} = 1$ and the bound (38) reduces to

$$\frac{3\mathcal{W}^2}{\epsilon(k+1)(k+2)} \leq \frac{\epsilon}{2}. \quad (40)$$

^cHar-Peled et al. [2007] use \mathbf{x}_i for the training data and \mathbf{w} for the hyperplane. Here we use \mathbf{z}_i and \mathbf{c} respectively to be consistent with our notation.

Solving for k shows that $\sqrt{6}\frac{\mathcal{W}}{\epsilon}$ iterations suffice to obtain an ϵ accurate solution of (39).

The Coreset SVM algorithm produces a multiplicative approximation. To compare with the bounds given by Har-Peled et al. [2007] we follow the same scheme described in Section 3.2. Let ρ^* denote the optimal margin, that is, $\min_{\mathbf{c} \in \mathbb{R}^d, \|\mathbf{c}\|=1} \max_i \langle \tilde{\mathbf{w}}_i, -\mathbf{c} \rangle$, and set $\epsilon = \rho^* \epsilon_M$. Substituting into (40) and solving for k shows that to produce a ϵ_M multiplicative approximation with our algorithm

$$\sqrt{6} \left(\frac{\mathcal{W}}{\rho^*} \right) \frac{1}{\epsilon_M} \leq k. \quad (41)$$

In contrast, Har-Peled et al. [2007] compute a Coreset SVM C of size

$$|C| = O \left(\left(\frac{\mathcal{W}}{\rho^*} \right)^2 \frac{1}{\epsilon_M} \right) \quad (42)$$

in $|C|$ iterations. Furthermore, the computational complexity of Coreset SVM is $O(md|C| + |C|T(|C|))$, where $T(|C|)$ is the cost of training a SVM on $|C|$ points. Since each iteration requires only $O(md)$ effort, our algorithm has an improved computational complexity of $O \left(md \left(\frac{\mathcal{W}}{\rho^*} \right) \frac{1}{\epsilon_M} \right)$. However, there is one notable difference between the two algorithms. Coreset SVM produces a sparse solution in terms of the number of support vectors, while our algorithm comes with no such guarantees.

5.2 Computing Polytope Distance

Given a set of points $S = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ the polytope distance is defined as the shortest distance ρ of any point inside the convex hull of S , $\text{conv}(S)$, to the origin [Gärtner and Jaggi, 2009]. Equivalently, we are looking for the vector with the smallest norm in $\text{conv}(S)$. A variant is to compute the distance between two polytopes given by the points $S_+ = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n\}$ and $S_- = \{\mathbf{z}'_1, \mathbf{z}'_2, \dots, \mathbf{z}'_{n'}\}$. This problem can be solved by finding the polytope distance of the Minkowski difference of $\text{conv}(S_+)$ and $\text{conv}(S_-)$ [Bennett and Bredensteiner, 1998]. Since the arguments for both cases are by and large very similar we will stick with the simpler formulation in this paper. The polytope distance problem has a number of applications in machine learning. A partial (and by no means exhaustive) list of relevant publications includes Gärtner and Jaggi [2009], Bennett and Bredensteiner [1998], and Keerthi et al. [2000].

To analyze this problem in our setting we start with a technical lemma. A version of this lemma also appears as Theorem A.2 in [Bennett and Bredensteiner \[1998\]](#).

Lemma 8 *The following two problems are duals of each other:*

$$\min_{\mathbf{c}} J(\mathbf{c}) := \max_i \langle \mathbf{c} - \mathbf{x}_i, \mathbf{c} \rangle \quad (43)$$

$$\max_{\mathbf{u} \in \Delta_n} D(\mathbf{u}) := -\frac{1}{4} \mathbf{u}^\top \mathbf{A} \mathbf{A}^\top \mathbf{u} = -\frac{1}{4} \left\| \mathbf{A}^\top \mathbf{u} \right\|^2, \quad (44)$$

where $\mathbf{A} = -[\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]^\top$.

Proof First rewrite the objective function in (43) as

$$J(\mathbf{c}) = \|\mathbf{c}\|^2 + \max_i \langle -\mathbf{x}_i, \mathbf{c} \rangle = \|\mathbf{c}\|^2 + \max_{\mathbf{u} \in \Delta_n} \langle \mathbf{u}, \mathbf{A} \mathbf{c} \rangle.$$

This can be identified with (4) by setting $Q_1 = \mathbb{R}^d$, $Q_2 = \Delta_n$, $f(\mathbf{c}) = \|\mathbf{c}\|^2$, and $g(\mathbf{u}) = 0$. The dual problem (44) can directly be read off from (7) by noting that the Fenchel dual of $\|\cdot\|^2$ is $\frac{1}{4} \|\cdot\|^2$. ■

Clearly (44) computes the vector with the smallest norm in $\text{conv}(S)$, which is equivalent to the polytope distance problem. Furthermore, (43) and (44) are identical to (12) and (13) respectively with $\mathbf{b} = \mathbf{0}$. Therefore the algorithm we described in Section 3 can be applied with minor modifications to yield a $O(nd/\sqrt{\epsilon})$ algorithm for this problem also. Since the [Gärtner and Jaggi \[2009\]](#) algorithm selects coresets, at most $O(1/\epsilon)$ components of \mathbf{u} are non-zero. However, in our case no such guarantees hold.

6 Conclusions and Future Work

We presented a new approximation algorithm for the MEB problem whose running time is $O(ndQ/\sqrt{\epsilon})$. Unlike existing algorithms, which rely heavily on geometric properties, our algorithm is motivated and derived purely from a convex optimization viewpoint. We extended our analysis to the MECF problem and obtain a $O(mndQ/\epsilon)$ algorithm. Not only does our treatment yield an algorithm with better bounds, but preliminary experimental results in Appendix D suggest that our algorithm is competitive on problems with a large number of data points.

A more general version of the MECF problem was studied by [Panigrahy \[2004\]](#). In his setting the convex polytope is allowed translations, magnifications, and rotations. A simple greedy approach (very reminiscent of coresets)

yields an $(1+\epsilon)$ multiplicative approximation algorithm which takes $O(1/\epsilon^2)$ iterations to converge. Please consult [Panigrahy \[2004\]](#) for details.

A natural question to ask is the following: Can our algorithms be extended to deal with arbitrary convex shapes? In other words, given an arbitrary convex shape can we find the optimal magnification and translation that is required to enclose the set of points $\{\mathbf{x}_j\}_{j=1}^n$ at hand. Unfortunately, a straightforward extension seems rather difficult. Even though it is well known that every convex shape can be described as an intersection of half planes, the number of half planes need not be finite. In such a case our MECP algorithm, which crucially relies on the number of half planes being finite, is clearly not applicable. Somewhat surprisingly, we are able to obtain a $O(1/\sqrt{\epsilon})$ algorithm for the MEB problem, even though a ball is made up of an intersection of infinitely many half planes. Clearly, the strong convexity of the objective function plays an important role in this context. We are currently trying to characterize such problems in the hope that this investigation will lead to efficient algorithms for a number of other related problems.

Rotations are a natural concept when working with geometric algorithms. This does not naturally carry over to our setting where we use convex optimization. In order to introduce rotations one has to work with orthogonal matrices, which significantly complicates the optimization strategy. A fruitful pursuit would be to investigate if the insights gained from coresets can be used to solve complicated optimization problems which involve orthogonal matrices efficiently.

Even though we are only beginning to scratch the surface on exploring connections between optimization and computational geometry, we firmly believe that this cross pollination will lead to exciting new algorithms in both areas.

References

- M. Badoiu and K.L. Clarkson. Optimal core-sets for balls. In *Computational Geometry: Theory and Applications*, 2002.
- K. P. Bennett and E. J. Bredensteiner. Geometry in learning. In C. Gorini, E. Hart, W. Meyer, and T. Phillips, editors, *Geometry at Work*, Washington, D.C., 1998. Mathematical Association of America. Available <http://www.math.rpi.edu/~bennek/geometry2.ps>.
- J. M. Borwein and A. S. Lewis. *Convex Analysis and Nonlinear Optimization: Theory and Examples*. CMS books in Mathematics. Canadian Mathematical Society, 2000.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, England, 2004.
- Kenneth L. Clarkson. Coresets, sparse greedy approximation, and the frank-wolfe algorithm. In *SODA '08: Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 922–931. Society for Industrial and Applied Mathematics, 2008.
- D. J. Elzinga and D. W. Hearn. The minimum covering sphere problem. *Management Science*, 19:96–104, 1972.
- Bernd Gärtner and Martin Jaggi. Coresets for polytope distance. In *Annual Symposium on Computational Geometry*, 2009.
- Sariel Har-Peled, Dan Roth, and Dav Zimak. Maximum margin coresets for active and noise tolerance learning. In *International Joint Conference on Artificial Intelligence*, pages 836–841, 2007.
- J.B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms, I and II*, volume 305 and 306. Springer-Verlag, 1993.
- S.S. Keerthi, S. K. Shevade, C. Bhattacharyya, and K. R. K. Murthy. A fast iterative nearest point algorithm for support vector machine classifier design. *IEEE Transactions on Neural Networks*, 11(1):124–136, January 2000.
- Nimrod Megiddo. Linear programming in linear time when the dimension is fixed. *J. ACM*, 31(1):114–127, 1984.

- Yurii Nesterov. Smooth minimization of non-smooth functions. *Math. Program.*, 103(1):127–152, 2005a.
- Yurii Nesterov. Excessive gap technique in nonsmooth convex minimization. *SIAM J. on Optimization*, 16(1):235–249, 2005b. ISSN 1052-6234.
- Yurii Nesterov. A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$. *Soviet Math. Docl.*, 269:543–547, 1983.
- Rina Panigrahy. Minimum enclosing polytope in high dimensions. *CoRR*, cs.CG/0407020, 2004.
- P. M. Pardalos and N. Kover. An algorithm for singly constrained class of quadratic programs subject to upper and lower bounds. *Mathematical Programming*, 46:321–328, 1990.
- B. Schölkopf and A. Smola. *Learning with Kernels*. MIT Press, Cambridge, MA, 2002.
- Ivor W. Tsang, James T. Kwok, and Pak-Ming Cheung. Core vector machines: Fast svm training on very large data sets. *J. Mach. Learn. Res.*, 6:363–392, 2005. ISSN 1532-4435.
- Ivor W. Tsang, András Kocsor, and James T. Kwok. Simpler core vector machines with enclosing balls. In *Proc. Intl. Conf. Machine Learning*, pages 911–918, 2007.
- Emo Welzl. Minimum enclosing disks (balls and ellipsoids). *Lecture Notes in Computer Science*, 555:359–370, 1991.
- E. Alper Yildirim. Two algorithms for the minimum enclosing ball problem. In *SIAM Journal on Optimization*, pages 1368–1391, 2008.

Appendix

A Proof of theorem 7

Our proof is by and large derived using results from [Nesterov \[2005b\]](#). We begin with a technical lemma.

Lemma 9 (*Lemma 7.2 of [Nesterov \[2005b\]](#)*) For any \mathbf{u} and $\bar{\mathbf{u}}$, we have

$$D(\mathbf{u}) + \langle \nabla D(\mathbf{u}), \bar{\mathbf{u}} - \mathbf{u} \rangle = \langle \mathbf{A}\mathbf{c}(\mathbf{u}) + \mathbf{b}, \bar{\mathbf{u}} \rangle + \|\mathbf{c}(\mathbf{u})\|^2.$$

Proof Direct calculation by using (13), (14), and (25) yields

$$\begin{aligned} D(\mathbf{u}) + \langle \nabla D(\mathbf{u}), \bar{\mathbf{u}} - \mathbf{u} \rangle &= \langle \mathbf{u}, \mathbf{b} \rangle - \frac{1}{4} \mathbf{u}^\top \mathbf{A} \mathbf{A}^\top \mathbf{u} + \left\langle \mathbf{b} - \frac{1}{2} \mathbf{A} \mathbf{A}^\top \mathbf{u}, \bar{\mathbf{u}} - \mathbf{u} \right\rangle \\ &= \langle \bar{\mathbf{u}}, \mathbf{b} \rangle - \left\langle \frac{1}{2} \mathbf{A} \mathbf{A}^\top \mathbf{u}, \bar{\mathbf{u}} \right\rangle + \frac{1}{4} \mathbf{u}^\top \mathbf{A} \mathbf{A}^\top \mathbf{u} \\ &= \langle \mathbf{A}\mathbf{c}(\mathbf{u}) + \mathbf{b}, \bar{\mathbf{u}} \rangle + \|\mathbf{c}(\mathbf{u})\|^2. \end{aligned}$$

■

We first show that the initial \mathbf{w}_1 and $\boldsymbol{\alpha}_1$ satisfy the excessive gap condition (9). Since $-D$ is L -l.c.g (from (15)), so

$$\begin{aligned} D(\mathbf{u}_1) &\geq D(\mathbf{u}_0) + \langle \nabla D(\mathbf{u}_0), \mathbf{u}_1 - \mathbf{u}_0 \rangle - \frac{L}{2} \|\mathbf{u}_1 - \mathbf{u}_0\|^2 \\ (\text{using defn. of } \mathbf{u}_1 \text{ and (19)}) &= \max_{\mathbf{u} \in \Delta_n} \left\{ D(\mathbf{u}_0) + \langle \nabla D(\mathbf{u}_0), \mathbf{u} - \mathbf{u}_0 \rangle - \frac{L}{2} \|\mathbf{u} - \mathbf{u}_0\|^2 \right\} \\ (\text{using lemma 9}) &= \max_{\mathbf{u} \in \Delta_n} \left\{ \langle \mathbf{u}, \mathbf{b} \rangle + \langle \mathbf{A}\mathbf{c}(\mathbf{u}_0), \mathbf{u} \rangle + \|\mathbf{c}(\mathbf{u}_0)\|^2 - \frac{L}{2} \|\mathbf{u} - \mathbf{u}_0\|^2 \right\} \\ (\text{using defn. of } d \text{ and } \mu_1) &= \max_{\mathbf{u} \in \Delta_n} \left\{ \langle \mathbf{u}, \mathbf{b} \rangle + \langle \mathbf{A}\mathbf{c}(\mathbf{u}_0), \mathbf{u} \rangle + \|\mathbf{c}(\mathbf{u}_0)\|^2 - \frac{\mu_1 \sigma}{2} \|\mathbf{u} - \mathbf{u}_0\|^2 \right\} \\ (\text{using } \mathbf{c}_1 = \mathbf{c}(\mathbf{u}_0)) &= \|\mathbf{c}_1\|^2 + \max_{\mathbf{u} \in \Delta_n} \left\{ \langle \mathbf{A}\mathbf{c}_1, \mathbf{u} \rangle + \langle \mathbf{u}, \mathbf{b} \rangle - \frac{\mu_1 \sigma}{2} \|\mathbf{u} - \mathbf{u}_0\|^2 \right\} \\ &\geq J_{\mu_1}(\mathbf{c}_1) \end{aligned}$$

which shows that our initialization indeed satisfies (9). Second, we prove by induction that the updates in Algorithm 1 maintain (9). We begin with two useful observations. Using (20) and the definition of τ_k , one can bound

$$\mu_{k+1} = (1 - \tau_k) \mu_k = \frac{6}{(k+3)(k+2)} \frac{L}{\sigma} \geq \tau_k^2 \frac{L}{\sigma}. \quad (45)$$

Let $\gamma := \mathbf{u}_{\mu_k}(\mathbf{c}_k)$. The optimality conditions for (18) imply $\langle \mu_k \sigma (\gamma - \mathbf{u}_0) - \mathbf{A} \mathbf{c}_k - \mathbf{b}, \mathbf{u} - \gamma \rangle \geq 0$ and hence

$$\mu_k \sigma \langle \gamma - \mathbf{u}_0, \mathbf{u} - \gamma \rangle \geq \langle \mathbf{A} \mathbf{c}_k + \mathbf{b}, \mathbf{u} - \gamma \rangle. \quad (46)$$

By using the update equation for \mathbf{c}_{k+1} and the convexity of $\|\cdot\|^2$

$$\begin{aligned} J_{\mu_{k+1}}(\mathbf{c}_{k+1}) &= \|\mathbf{c}_{k+1}\|^2 + \max_{\mathbf{u} \in \Delta_n} \left\{ \langle \mathbf{A} \mathbf{c}_{k+1}, \mathbf{u} \rangle + \langle \mathbf{u}, \mathbf{b} \rangle - \frac{\mu_{k+1} \sigma}{2} \|\mathbf{u} - \mathbf{u}_0\|^2 \right\} \\ &= \|(1 - \tau_k) \mathbf{c}_k + \tau_k \mathbf{c}(\beta_k)\|^2 \\ &\quad + \max_{\mathbf{u} \in \Delta_n} \left\{ (1 - \tau_k) \langle \mathbf{A} \mathbf{c}_k, \mathbf{u} \rangle + \tau_k \langle \mathbf{A} \mathbf{c}(\beta_k), \mathbf{u} \rangle + \langle \mathbf{u}, \mathbf{b} \rangle - (1 - \tau_k) \frac{\mu_k \sigma}{2} \|\mathbf{u} - \mathbf{u}_0\|^2 \right\} \\ &\leq \max_{\mathbf{u} \in \Delta_n} \{ (1 - \tau_k) T_1 + \tau_k T_2 \}, \end{aligned}$$

where

$$\begin{aligned} T_1 &= \left[-\frac{\mu_k \sigma}{2} \|\mathbf{u} - \mathbf{u}_0\|^2 + \langle \mathbf{A} \mathbf{c}_k + \mathbf{b}, \mathbf{u} \rangle + \|\mathbf{c}_k\|^2 \right] \quad \text{and} \\ T_2 &= \left[\langle \mathbf{A} \mathbf{c}(\beta_k) + \mathbf{b}, \mathbf{u} \rangle + \|\mathbf{c}(\beta_k)\|^2 \right]. \end{aligned}$$

T_1 can be bounded as follows

$$\begin{aligned} T_1 &= -\frac{\mu_k \sigma}{2} \|\mathbf{u} - \mathbf{u}_0\|^2 + \langle \mathbf{A} \mathbf{c}_k + \mathbf{b}, \mathbf{u} \rangle + \|\mathbf{c}_k\|^2 \\ &= -\frac{\mu_k \sigma}{2} \|\mathbf{u} - \gamma\|^2 - \frac{\mu_k \sigma}{2} \|\gamma - \mathbf{u}_0\|^2 - \mu_k \sigma \langle \gamma - \mathbf{u}_0, \mathbf{u} - \gamma \rangle \\ &\quad + \langle \mathbf{A} \mathbf{c}_k + \mathbf{b}, \mathbf{u} \rangle + \|\mathbf{c}_k\|^2 \\ (\text{using (46)}) &\leq -\frac{\mu_k \sigma}{2} \|\mathbf{u} - \gamma\|^2 - \frac{\mu_k \sigma}{2} \|\gamma - \mathbf{u}_0\|^2 - \langle \mathbf{A} \mathbf{c}_k + \mathbf{b}, \mathbf{u} - \gamma \rangle \\ &\quad + \langle \mathbf{A} \mathbf{c}_k + \mathbf{b}, \mathbf{u} \rangle + \|\mathbf{c}_k\|^2 \\ &= -\frac{\mu_k \sigma}{2} \|\mathbf{u} - \gamma\|^2 - \frac{\mu_k \sigma}{2} \|\gamma - \mathbf{u}_0\|^2 + \langle \mathbf{A} \mathbf{c}_k + \mathbf{b}, \gamma \rangle + \|\mathbf{c}_k\|^2 \\ (\text{using defn. of } \gamma) &= -\frac{\mu_k \sigma}{2} \|\mathbf{u} - \gamma\|^2 + J_{\mu_k}(\mathbf{w}_k) \\ (\text{using induction assumption}) &\leq -\frac{\mu_k \sigma}{2} \|\mathbf{u} - \gamma\|^2 + D(\mathbf{u}_k) \\ (\text{using concavity of } D) &\leq -\frac{\mu_k \sigma}{2} \|\mathbf{u} - \gamma\|^2 + D(\beta_k) + \langle \nabla D(\beta_k), \mathbf{u}_k - \beta_k \rangle, \end{aligned}$$

while T_2 can be simplified by using Lemma 9:

$$T_2 = \langle \mathbf{u}, \mathbf{b} \rangle + \langle \mathbf{A} \mathbf{c}(\beta_k), \mathbf{u} \rangle + \|\mathbf{c}(\beta_k)\|^2 = D(\beta_k) + \langle \nabla D(\beta_k), \mathbf{u} - \beta_k \rangle.$$

Putting the upper bounds on T_1 and T_2 together, and using (45) we obtain the following result.

$$\begin{aligned}
J_{\mu_{k+1}}(\mathbf{c}_{k+1}) &\leq \max_{\mathbf{u} \in \Delta_n} \left\{ (1 - \tau_k) \left[-\frac{\mu_k \sigma}{2} \|\mathbf{u} - \boldsymbol{\gamma}\|^2 + D(\boldsymbol{\beta}_k) + \langle \nabla D(\boldsymbol{\beta}_k), \mathbf{u}_k - \boldsymbol{\beta}_k \rangle \right] \right. \\
&\quad \left. + \tau_k [D(\boldsymbol{\beta}_k) + \langle \nabla D(\boldsymbol{\beta}_k), \mathbf{u} - \boldsymbol{\beta}_k \rangle] \right\} \\
&\leq D(\boldsymbol{\beta}_k) + \max_{\mathbf{u} \in \Delta_n} \left\{ -\tau_k^2 \frac{L}{2} \|\mathbf{u} - \boldsymbol{\gamma}\|^2 + \langle \nabla D(\boldsymbol{\beta}_k), (1 - \tau_k)\mathbf{u}_k + \tau_k \mathbf{u} - \boldsymbol{\beta}_k \rangle \right\}.
\end{aligned} \tag{47}$$

Let $\mathbf{v} = (1 - \tau_k)\mathbf{u}_k + \tau_k \mathbf{u}$. By using the definition of $\boldsymbol{\beta}_k$ from Algorithm 1 observe that

$$(1 - \tau_k)\mathbf{u}_k + \tau_k \mathbf{u} - \boldsymbol{\beta}_k = \tau_k(\mathbf{u} - \boldsymbol{\gamma}) = \mathbf{v} - \boldsymbol{\beta}_k. \tag{48}$$

Furthermore, $\mathbf{v} \in \Delta_n$ since it is a convex combination of $\mathbf{u}_k \in \Delta_n$ and $\mathbf{u} \in \Delta_n$. Plugging (48) into (47)

$$J_{\mu_{k+1}}(\mathbf{c}_{k+1}) \leq D(\boldsymbol{\beta}_k) + \max_{\mathbf{v} \in \Delta_n} \left\{ -\frac{L}{2} \|\mathbf{v} - \boldsymbol{\beta}_k\|^2 + \langle \nabla D(\boldsymbol{\beta}_k), \mathbf{v} - \boldsymbol{\beta}_k \rangle \right\}$$

$$\begin{aligned}
&\text{(using (19) and defn. of } \mathbf{u}_{k+1}) = D(\boldsymbol{\beta}_k) + \langle \nabla D(\boldsymbol{\beta}_k), \mathbf{u}_{k+1} - \boldsymbol{\beta}_k \rangle - \frac{L}{2} \|\mathbf{u}_{k+1} - \boldsymbol{\beta}_k\|^2 \\
&\text{(Since } -D \text{ is } L\text{-l.c.g.)} \leq D(\mathbf{u}_{k+1}).
\end{aligned}$$

B A linear time algorithm for a box constrained diagonal QP with a single linear equality constraint

In this section, we focus on the following simple QP:

$$\begin{aligned}
&\min \frac{1}{2} \sum_{i=1}^n d_i^2 (\alpha_i - m_i)^2 \\
&s.t. \quad l_i \leq \alpha_i \leq u_i \quad \forall i \in [n]; \\
&\quad \sum_{i=1}^n \sigma_i \alpha_i = z.
\end{aligned} \tag{49}$$

Without loss of generality, we assume $l_i < u_i$ and $d_i \neq 0$ for all i . Also assume $\sigma_i \neq 0$ because otherwise α_i can be solved independently. To make the feasible region nonempty, we also assume

$$\sum_i \sigma_i (\delta(\sigma_i > 0) l_i + \delta(\sigma_i < 0) u_i) \leq z \leq \sum_i \sigma_i (\delta(\sigma_i > 0) u_i + \delta(\sigma_i < 0) l_i).$$

The algorithm we describe below stems from [Pardalos and Kuvshinov \[1990\]](#) and finds the exact optimal solution in $O(n)$ time.

With a simple change of variable $\beta_i = \sigma_i(\alpha_i - m_i)$, the problem is simplified as

$$\begin{aligned}
\min \quad & \frac{1}{2} \sum_{i=1}^n \bar{d}_i^2 \beta_i^2 \\
s.t. \quad & l'_i \leq \beta_i \leq u'_i \quad \forall i \in [n]; \quad \text{where} \\
& \sum_{i=1}^n \beta_i = z', \\
& l'_i = \begin{cases} \sigma_i(l_i - m_i) & \text{if } \sigma_i > 0 \\ \sigma_i(u_i - m_i) & \text{if } \sigma_i < 0 \end{cases}, \\
& u'_i = \begin{cases} \sigma_i(u_i - m_i) & \text{if } \sigma_i > 0 \\ \sigma_i(l_i - m_i) & \text{if } \sigma_i < 0 \end{cases}, \\
& \bar{d}_i^2 = \frac{d_i^2}{\sigma_i^2}, \quad z' = z - \sum_i \sigma_i m_i.
\end{aligned}$$

We derive its dual via the standard Lagrangian.

$$L = \frac{1}{2} \sum_i \bar{d}_i^2 \beta_i^2 - \sum_i \rho_i^+ (\beta_i - l'_i) + \sum_i \rho_i^- (\beta_i - u'_i) - \lambda \left(\sum_i \beta_i - z' \right).$$

Taking derivative:

$$\frac{\partial L}{\partial \beta_i} = \bar{d}_i^2 \beta_i - \rho_i^+ + \rho_i^- - \lambda = 0 \quad \Rightarrow \quad \beta_i = \bar{d}_i^{-2} (\rho_i^+ - \rho_i^- + \lambda). \quad (50)$$

Substituting into L , we get the dual optimization problem

$$\begin{aligned}
\min D(\lambda, \rho_i^+, \rho_i^-) &= \frac{1}{2} \sum_i \bar{d}_i^{-2} (\rho_i^+ - \rho_i^- + \lambda)^2 - \sum_i \rho_i^+ l'_i + \sum_i \rho_i^- u'_i - \lambda z' \\
s.t. \quad & \rho_i^+ \geq 0, \quad \rho_i^- \geq 0 \quad \forall i \in [n].
\end{aligned}$$

Taking derivative of D with respect to λ , we get:

$$\sum_i \bar{d}_i^{-2} (\rho_i^+ - \rho_i^- + \lambda) - z' = 0. \quad (51)$$

The KKT condition gives:

$$\rho_i^+(\beta_i - l'_i) = 0, \quad (52a)$$

$$\rho_i^-(\beta_i - u'_i) = 0. \quad (52b)$$

Now we enumerate four cases.

1. $\rho_i^+ > 0, \rho_i^- > 0$. This implies that $l'_i = \beta_i = u'_i$, which is contradictory to our assumption.

2. $\rho_i^+ = 0, \rho_i^- = 0$. Then by (50), $\beta_i = \bar{d}_i^{-2}\lambda \in [l'_i, u'_i]$, hence $\lambda \in [\bar{d}_i^2 l'_i, \bar{d}_i^2 u'_i]$.

3. $\rho_i^+ > 0, \rho_i^- = 0$. Now by (52) and (50), we have $l'_i = \beta_i = \bar{d}_i^{-2}(\rho_i^+ + \lambda) > \bar{d}_i^{-2}\lambda$, hence $\lambda < \bar{d}_i^2 l'_i$ and $\rho_i^+ = \bar{d}_i^2 l'_i - \lambda$.

4. $\rho_i^+ = 0, \rho_i^- > 0$. Now by (52) and (50), we have $u'_i = \beta_i = \bar{d}_i^{-2}(-\rho_i^- + \lambda) < \bar{d}_i^{-2}\lambda$, hence $\lambda > \bar{d}_i^2 u'_i$ and $\rho_i^- = -\bar{d}_i^2 u'_i + \lambda$.

In sum, we have $\rho_i^+ = [\bar{d}_i^2 l'_i - \lambda]_+$ and $\rho_i^- = [\lambda - \bar{d}_i^2 u'_i]_+$. Now (51) turns into

$$f(\lambda) := \sum_i \underbrace{\bar{d}_i^{-2}([\bar{d}_i^2 l'_i - \lambda]_+ - [\lambda - \bar{d}_i^2 u'_i]_+ + \lambda)}_{=: h_i(\lambda)} - z' = 0. \quad (53)$$

In other words, we only need to find the root of $f(\lambda)$ in (53). $h_i(\lambda)$ is given by

$$h_i(\lambda) = \begin{cases} l'_i & \text{if } \lambda < \bar{d}_i^2 l'_i \\ \lambda \bar{d}_i^{-2} & \text{if } \bar{d}_i^2 l'_i \leq \lambda \leq \bar{d}_i^2 u'_i \\ u'_i & \text{if } \lambda > \bar{d}_i^2 u'_i \end{cases} \quad (54)$$

Note that $h_i(\lambda)$ is a monotonically increasing function of λ , so the whole $f(\lambda)$ is monotonically increasing in λ . Since $f(\infty) \geq 0$ by $z' \leq \sum_i u'_i$ and $f(-\infty) \leq 0$ by $z' \geq \sum_i l'_i$, the root must exist. Considering that f has at most $2n$ kinks (nonsmooth points) and is linear between two adjacent kinks, the simplest idea is to sort $\{\bar{d}_i^2 l'_i, \bar{d}_i^2 u'_i : i \in [n]\}$ into $s^{(1)} \leq \dots \leq s^{(2n)}$. If $f(s^{(i)})$ and $f(s^{(i+1)})$ have different signs, then the root must lie between them and can be easily found because f is linear in $[s^{(i)}, s^{(i+1)}]$. This algorithm takes at least $O(n \log n)$ time because of sorting.

Algorithm 2: $O(n)$ algorithm to find the root of $f(\lambda)$. Ignoring boundary condition checks.

Input: Function f

Output: λ^* : Root of f

```

1 Initialize: Set kink set  $S \leftarrow \{\bar{d}_i^2 l'_i : i \in [n]\} \cup \{\bar{d}_i^2 u'_i : i \in [n]\}.$ ;
2 while  $|S| > 2$  do
3   Find median of  $S$ :  $m \leftarrow \text{MED}(S)$ .
4   if  $f(m) \geq 0$  then
5      $S \leftarrow \{x \in S : x \leq m\}.$ 
6   else
7      $S \leftarrow \{x \in S : x \geq m\}.$ 
8 Return root  $\frac{lf(u)-uf(l)}{f(u)-f(l)}$  where  $S = \{l, u\}.$ 

```

However, this complexity can be reduced to $O(n)$ by making use of the fact that the median of n (unsorted) elements can be found in $O(n)$ time. Notice that due to the monotonicity of f , the median of a set S gives exactly the median of function values, *i.e.*, $f(\text{MED}(S)) = \text{MED}(\{f(x) : x \in S\})$. Algorithm 2 sketches the idea of binary search. The while loop terminates in $\log_2(2n)$ iterations because the set S is halved in each iteration. And in each iteration, the time complexity is linear to $|S|$, the size of current S . So the total complexity is $O(n)$. Note the evaluation of $f(m)$ potentially involves summing up n terms as in (53). However by some clever aggregation of slope and offset, this can be reduced to $O(|S|)$.

C Learning SVMs with MEB

Kernel methods in general and support vector machines (SVMs) in particular have received significant recent research interest in machine learning [Schölkopf and Smola, 2002]. Underlying a SVM is a simple geometric idea. Given a training set $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ of n points \mathbf{x}_i labeled by $y_i \in \{\pm 1\}$ the aim is to find the hyperplane which maximizes the margin of separation between points from different classes. This is compactly written as the following

optimization problem (see [Schölkopf and Smola \[2002\]](#) for details):

$$\min_{\mathbf{w} \in \mathbb{R}^d} \quad \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^n \xi_i \quad (55a)$$

$$\text{s.t.} \quad y_i(\langle \mathbf{w}, \mathbf{x}_i \rangle + b) \geq 1 - \xi_i \text{ for all } i. \quad (55b)$$

Standard duality arguments (see *e.g.* [Boyd and Vandenberghe \[2004\]](#)) yield the following dual Quadratic Programming (QP) problem

$$\max_{\boldsymbol{\alpha} \in \mathbb{R}^m} \quad \langle \boldsymbol{\alpha}, \mathbf{e} \rangle - \frac{1}{2} \boldsymbol{\alpha}^\top K \boldsymbol{\alpha} \quad (56a)$$

$$\text{s.t.} \quad \sum_{i=1}^n \alpha_i y_i = 0 \quad (56b)$$

$$0 \leq \alpha_i \leq C \text{ for all } i. \quad (56c)$$

Here K is a $m \times m$ matrix whose entries are given by $y_i y_j \langle \mathbf{x}_i, \mathbf{x}_j \rangle$ and \mathbf{e} denotes the vector of all ones. Since the dual only depends on \mathbf{x} via the inner products $\langle \mathbf{x}_i, \mathbf{x}_j \rangle$ one can employ the kernel trick: map ξ_i into a feature space via $\phi(\mathbf{x}_i)$ and compute the dot product in the feature space by using a kernel function $k(\mathbf{x}_i, \mathbf{x}_j) := \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle$ [Schölkopf and Smola \[2002\]](#). The kernel trick makes SVM rather powerful because simple linear decision boundaries in feature space map into non-linear decision boundaries in the original space where the datapoints live.

A number of different techniques have been proposed for solving the quadratic problem associated with SVMs. Of particular interest in our context is the Core Vector Machine (CVM) [[Tsang et al., 2007, 2005](#)]. The key idea of the CVM is the observation that solving (56) is equivalent to finding the MEB of the feature vectors. The MEB problem in feature space can be written as (see [Tsang et al. \[2007\]](#) for details)

$$\min_{\mathbf{c} \in \mathbb{R}^d, R \in \mathbb{R}} \quad R^2 \quad (57a)$$

$$\|\mathbf{c} - \phi(\mathbf{x}_i)\|^2 \leq R^2 \text{ for all } i. \quad (57b)$$

The dual of the above minimization problem then becomes (also see (13))

$$\max_{\boldsymbol{\alpha}} \quad \sum_{i=1}^n \alpha_i K_{i,i} - \boldsymbol{\alpha}^\top K \boldsymbol{\alpha} \quad (58a)$$

$$\text{s.t.} \quad \alpha_i \geq 0, \quad \sum_i \alpha_i = 1 \quad (58b)$$

where $K_{i,j} = \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle$ as before. In particular, if each $K_{i,i}$ equals a constant c then it can be shown that by a simple transformation that the standard SVM dual (56) and the CVM dual (58) can be identified. Therefore, every iteration of the CVM algorithm [Tsang et al., 2005] identifies an active set of points, and computes the MEB of this active set. This is done via the coreset algorithm of Panigrahy [2004], hence the name core vector machine. In fact, our algorithm for the MEB can directly be plugged into the CVM, and improves the rates of convergence of the inner iteration from $O(nd/\epsilon)$ to $O(nd/\sqrt{\epsilon})$. We hope that our algorithm with improved convergence rates can also be used in similar machine learning algorithms to speed up their convergence.

D Experimental Results

The aim of our experiments is to demonstrate the efficacy of Algorithm 1 and compare its performance with existing MEB algorithms in terms of running time and number of iterations required for convergence. Following [Yildirim, 2008] we generate data using a random multivariate Gaussian distribution and vary n the number of data points and d the dimensions. For each fixed n and d we generate 5 random datasets and report the average performance of our algorithm with the multiplicative guarantee $\epsilon = 10^{-3}$ in Table 1. Recall from Section 3.2 that the multiplicative guarantee ϵ is equivalent to the additive tolerance $\epsilon\mathcal{P}^2$, where \mathcal{P} is chosen via

$$\mathcal{P} = \frac{1}{2} \max_{\mathbf{x}_i, \mathbf{x}_j} \|\mathbf{x}_i - \mathbf{x}_j\|.$$

For reference we also reproduce the results reported in Table 1 of [Yildirim, 2008]. Although not a fair comparison, the CPU times gives an indication of the relative performance of various algorithms. In the table BC refers to the coreset algorithm of Badoiu and Clarkson [2002] while A1 and A2 are MEB algorithms implemented in [Yildirim, 2008].

Our algorithm performs significantly better than BC and A1, while being comparable to A2. Furthermore, our algorithm usually takes smaller number of iterations and particularly shines when the number of points is large. Our implementation is preliminary, and we believe that our algorithm will benefit from practical speed ups in much the same way that algorithm A2 was obtained by improvising A1 to get rid of redundancies.

n	d	Time				Iterations			
		A1	A2	BC	Ours	A1	A2	BC	Ours
500	10	0.06	0.03	0.12	0.04	168.7	44.5	435.5	44.2
1000	10	0.15	0.03	0.14	0.10	330.7	41.6	344.4	54.5
5000	20	1.7	0.36	3.11	1.08	246.8	46	464.2	69.7
10000	20	4.46	0.58	4.65	3.40	319.2	36.3	334.4	105.2
30000	30	27	6.45	24.59	10.43	446.4	103.6	409	77.8
50000	50	71.62	16.87	68.78	18.69	429.8	98.4	415.1	54.5
100000	100	287.99	77.74	268.11	83.18	451.7	119	422.6	63

Table 1: Computational Results with $n \gg d$ ($\epsilon = 10^{-3}$)